Many of the LANs operated by large businesses operate according to IEEE 802.1 standards. These standards provide protocols that enable the businesses to partition their LANs into multiple virtual LANS (VLANs). So, for example, a large business may partition its LANs into separate VLANs for different departments or operations of the business, like Finance, Manufacturing, Design and Legal. The Walker et al (US 5,862,338) and Yuasa et al (US 6,085,238) references cited by the Examiner describe IEEE 802.1 LANs having VLAN capabilities as required by that standard.

The Applicants invention goes beyond the VLAN capabilities required by IEEE 802.1 to enable Carrier Network Service Providers to provide a very large number of VLANs on shared network facilities in a manner which makes more efficient use of the Carrier Network Service Provider network than the typical leased dedicated circuits, while preserving isolation between the data communications of separate customers, and while preserving the VLAN partitioning of those customers within their own data networks.

The Applicants claim communications networks, routing devices for such communications networks and methods of routing packets in which each packet entering a network at an ingress virtual port is assigned a respective egress address and routed according to that respective egress address. The assigned egress address corresponds to a respective destination address of the entering packet when a correspondence between the destination address and an egress address is known. When no correspondence between the destination address and an egress address is known, the assigned address is a broadcast egress address which corresponds to a distinct set of virtual ports, the distinct set of virtual ports comprising the ingress virtual port. Consequently, broadcasting of the entering packet is restricted to the distinct set of virtual ports that includes the ingress virtual port.

The Examiner rejected claims 1-18, 21-32 and 45-50 under 35 USC 103(a) as being unpatentable over Walker et al (US 5,862,338).

Walker et al discloses a network switch that operates according to the IEEE 802.1 standard with regard to VLAN functionality. The IEEE 802.1 standard requires that a header of each frame of data carry a VLAN tag that identifies the VLAN for which the data frame is intended. When no correspondence between

 incoming packet is known by the switch, the switch

reads the header of the incoming packet and broadcasts the data frames based on the VLAN tag to only those ports that, according to routing tables stored at the switches, are participating in that VLAN.

Such IEEE 802.1 VLANs are described at page 1, line 31 to page 2, line 11 of the Applicants' specification. As noted there, the 12 bit capacity of the VLAN tag specified by the IEEE 802.1 standard limits the number of distinct VLANs to 4095. Carrier Network Service Providers need to support many more than 4095 distinct customers on a shared network. Moreover, many customers of the Service Providers are already using the VLAN tag to partition their own networks and do not want Carrier Network Service Providers to disrupt such partitioning by changing the VLAN identifiers on packets traversing the Carrier Networks.

Furthermore, Walker et al uses VLAN tags to route only those packets having a destination MAC address which has not yet been captured in routing tables of the switch of Walker et al. This implies that if multiple switches of the type taught by Walker et al were connected to form a large network, all of the switches would need to populate their routing tables as described at Figure 12 and at column 60, lines 21-31 of Walker et al requiring considerable processing effort and Walker's routing tables would need to be large enough to accommodate a very large number of destination addresses. This further implies that each device connected to a network of switches according to Walker et al must have a unique MAC address.

The Applicants avoid such limitations of IEEE 802.1 VLANs by using assigned egress addresses rather than destination MAC addresses and VLAN tags to route broadcast packets. Because the number of different broadcast addresses is much greater than the number of different VLAN identifiers permitted under the IEEE 802.1 standard, the Applicants can provide a larger number of isolated virtual private networks than can a standard IEEE 802.1 VLAN network. Moreover, because the Applicants do not disturb VLAN tags that may be used by their corporate customers, the partitioning of each corporate customer's VLAN remains intact.

Furthermore, by using assigned egress addresses rather than incoming destination addresses to route packets within the network, the routing tables of intermediate routing devices 16 of the Applicants' network need only

ted with routing instructions for the number of egress

addresses corresponding to the number of virtual ports on the network, which can be considerably lower than the number of destination MAC addresses connected to the network for a Service Provider network. This considerably reduces the size of these tables and the amount of packet processing required to populate them and keep them up to date. Moreover, shorter routing tables may enable faster routing of packets at the intermediate routing devices. Because egress addresses are assigned based on both the ingress virtual port and the destination MAC address, each source and destination MAC addresses need only be unique to its ingress virtual port, not to the entire network. Consequently, Service Providers can permit their corporate customers to administer their own MAC address space without having to ensure that locally administered MAC addresses used by one corporate customer do not overlap those used by another corporate customer.

The Examiner relied on column 61, lines 55-57; and column 63, lines 14-23 of Walker et al to show that "Each packet is assigned ... a broadcast address when no correspondence between the destination address and an egress address is known". However, at column 63, lines 14-18 Walker et al clearly state that the VLAN tag is used to route broadcast packets, not an assigned egress address.

In summary, Walker et al neither teach nor suggest assignment of a respective egress address to each packet entering a network and routing the packet according to the egress address (as opposed to the destination address that the packet already carries as it enters the network) to enable more efficient routing of the packet across the network. Moreover, Walker et al neither teach nor suggest assigning a broadcast address corresponding to a distinct set of virtual ports including the ingress virtual port (as opposed to a VLAN tag according the IEEE 802.1 standard) so as to increase the number of distinct VLANs that can be supported by a carrier without disrupting the partitioning of VLANs of corporate customers of a Carrier Service Provider. Consequently, the Applicants submit that claims 1-18, 21-32 and 45-50 are patentably distinct from Walker et al. The Applicants therefore request withdrawal of the rejection of these claims under 35 USC 103(a).

The Applicants' dependent claims contain advantageous features that further distinguish inventions defined by these claims from the disclosure of Walker et al.

The Applicants note that claims 6, 8, 16, 18, 26, 28 and 32 further require the assignment of a respective ingress address to each packet entering the network and use of the assigned ingress address to populate address association tables. The Examiner refers to Figure 9C of Walker et al to show an address table based on source and destination MAC addresses, and refers to Figure 12, steps 1218, 1222, 1242 and 1250 together with column 63, lines 14-23 to show population of the address tables. The Applicants note that Walker et al do not assign ingress addresses to packets entering the network. Moreover, Walker et al populate routing tables using source MAC addresses that packets already carry when they enter the network, not ingress addresses assigned to the packets as they enter the network. The use of assigned ingress addresses enables the Applicants to use a smaller address space containing based on the addresses of virtual ports on the network rather than the MAC addresses of terminals connected to the network – and this smaller address space permits smaller address tables which require less processing to populate and maintain.

The Applicants note that claims 7-8, 17-18 and 27-28 further require that the respective assigned egress address be added to each packet entering the network and used for routing the packet through the network, and that the respective assigned egress address be removed from each packet leaving the network. This encapsulation of each packet as it enters the network and decapsulation of each packet as it leaves the network leaves the customers' packets intact, a clear benefit to customers of the carrier VLAN service.

The Examiner admitted that Walker does not disclose adding to each packet the egress address assigned to the packet when the packet enters the network and removing the assigned egress address when the packet leaves the network. However, the Examiner argued that "One skilled in the art would recognize that switches typically use a temporary internal header for processing packets with the switch (i.e. route to a destination port, queue by priority, schedule according to time or bandwidth, etc.) and once the packet it routed to a destination port, the internal header is removed ... Therefore it would have been obvious to one skilled in the art at the time the invention was made to use an internal header in the invention of Walker as a matter of design choice."

The Applicants submit that the use of a temporary internal header within

from adding an assigned egress address to a packet

when the packet enters the network, using that assigned egress address to route the packet at any intervening switches or routing devices at multiple nodes within the network, and removing the assigned address from the packet when the packet is leaving the network. In particular, the added egress address required by Applicants' claims 7-8, 17-18 and 27-28 is not used only for processing of the packet within a single switch and stripped from the packet before it leaves the switch. Consequently, the added address can be used to route the packet at other switches or routing devices throughout the network, not just in the switch where the packet enters the network. This enables efficient routing of the packets through the network without excessive delay.

The Applicants further note that claims 9-10 and 29-30 require that packets be routed via a restricted set of trunks containing only those trunks required to reach virtual ports in the distinct set of virtual ports corresponding to assigned broadcast or multicast addresses. No such use of a restricted set of trunks is taught or suggested by Walker et al. The Examiner relied on ThunderLAN Port Interfaces (TPIs) 202 shown in Figure 2 to show network trunks. Comparing Figure 2 to Figure 1, it is clear that the TPIs are interfaces to LAN segments, not network trunks. Moreover, there is no suggestion in Walker et al that packets be routed to restricted sets of trunks.

The Applicants further note that claim 31 requires routers having IEEE 802.1 switching functionality adapted to packets encapsulated with ingress and egress addresses. The Examiner notes that the network of Walker et al is Ethernet-based, referring to column 7, line 30. However, the Applicants note that Walker et al neither teaches nor suggests routers that are adapted to packets encapsulated with ingress and egress addresses.

The Applicants further note that claim 32 requires a respective address assigner for each distinct subset of virtual ports. The Examiner refers to Figures 9C and 12 and column 63, lines 14-23 in discussing claim 32, but none of these teach or suggest a respective address assigner for each of plural distinct subsets of ports. In fact, the discussion of Figure 9C in column 50 and the discussion of Figure 12 in columns 57+ relate to routing of packets based on their address and VLAN identifier fields – these figures do not teach or suggest the assignment of addresses to packets. Using separate address assigners for each distinct subset

 and breaches of security if multiple customers use

overlapping locally-administered MAC address spaces and can enhance security since each address assigner can be programmed to assign completely distinct sets of egress addresses.

In view of the above arguments, the Applicants submit that claims 1-18, 21-32 and 45-50 are patentable over Walker et al. The Applicants therefore request withdrawal of the rejection of these claims under 35 USC 103 and allowance of these claims.

The Examiner rejected claims 19-20 under 35 USC 103(a) as being unpatentable over Walker et al and Yuasa et al (US 6,085,238). The Examiner stated that Walker et al does not disclose trunks interconnecting routers of a network. However the Examiner argued that Yuasa et al does disclose trunks interconnecting routers of a network, so it would be obvious to a skilled person to apply the teachings of Walker et al to a network comprising routers connected by trunks, as disclosed by Yuasa et al to provide connectivity between VLAN members that are too far apart to be connected to the same switch or node.

The Applicants note that, like Walker et al, Yuasa et al neither teach nor suggest assignment of an egress address to a packet entering a network via an ingress virtual port, the egress address corresponding to a destination address of the entering packet when a correspondence between the destination address and an egress address is known, and the respective egress address being a broadcast address corresponding to a distinct set of virtual ports comprising the ingress virtual port when no correspondence between the destination address and an egress address is known. Like Walker et al, Yuasa et al route packets based on a destination MAC address that the packet carries when it enters the network or, when no routing for the destination MAC address is known, based on a standard IEEE 802.1 VLAN tag rather than an assigned egress address, incurring the limitations of discussed above with reference to Walker et al. Consequently, Yuasa et al does not address the deficiencies of Walker et al with respect to the Applicants' independent claim 11 (from which claims 19 and 20 depend).

In fact, Yuasa et al teach directly away from some of the Applicants claims.

In particular, at column 7, lines 21 to 27 and elsewhere, Yuasa et al teach that virtual group membership of packets entering a network should be

 destination addresses of the packet, not from the port

on which the packet enters the network. Yuasa et al do not rely on the identity of the ingress port to identify the virtual group membership because Yuasa et al wish to accommodate frequent terminal moves by enabling their network to maintain the virtual group membership of terminals even when terminals are moved from one network port to another. This may be acceptable for VLANs that are confined to a single business. However, this would clearly be unacceptable for VLAN services provided by Carrier Service Providers since it enables terminals to declare their virtual group membership and this is clearly insecure when the terminals may belong to distinct Carrier Service Provider customers who must not be permitted to access one another's communications.

Moreover, at column 7, lines 21 to 27 and elsewhere, Yuasa et al suggest that no extra headers or tags should be added to packets to ensure compatibility in multi-vendor networks in contrast to certain of Applicants' claims (e.g. claim 17) that require encapsulation of packets entering a network.

Furthermore, Applicants' claims 19 and 20 require that packets be routed via a restricted set of trunks containing only those trunks required to reach virtual ports in the distinct set of virtual ports corresponding to assigned broadcast or multicast addresses. No such use of a restricted set of trunks is taught or suggested by Yuasa et al. The Examiner relied on Figure 4 and column 21, lines 52-65 with respect to claims 19 and 20, but there is no suggestion here that packets be routed to restricted sets of trunks.

In view of the above arguments, the Applicants submit that claims 19-20 are patentable over Walker et al and Yuasa et al. The Applicants therefore request withdrawal of the rejection of these claims under 35 USC 103 and allowance of these claims.

The Examiner objected to claims 33-44 as being dependent upon a rejected base claim, but indicated that these claims would be allowable if rewritten in independent form including all limitations of the base claim. The Applicants have amended claims 32 and 41 to put these claims in independent form. Claims 33-40 and 42-44 depend from these claims. While the Applicants submit that these claims are patentable without these amendments based on the arguments detailed above, the Applicants have made these amendments to expedite allowance of these claims.

The Applicants have rewritten claim 32 in independent form rather than claims 33 and 35 to provide a simpler amendment. The Applicants argued above that claim 32 is clearly patentable since its features are neither taught nor suggested by the cited references. Moreover, claims 33-40, which depend from claim 32, have further features such that the Examiner considered these claims to be allowable if they were made dependent on an allowable base claim.)

The Commissioner is hereby authorized to charge any additional fee which may be required, or credit any overpayment to Deposit Account #14-1315.

The Applicants request allowance of the amended application based on the discussion above.

Respectfully submitted,

By: _____

C.W. Junkin
Reg. No. 32,812

CWJ:jp

c/o NORTEL NETWORKS LIMITED
IP Law Group, P.O. Box 3511,
Station "C", Ottawa, Ontario K1Y 4H7
Canada

Phone: (613) 768-3013
Date: December 20, 2002

CERTIFICATE OF FACSIMILE TRANSMISSION

I hereby certify that this paper (1 page) is being facsimile transmitted to the Patent and Trademark Office on the date shown below.

_____

Judy Potvin, (613) 768-3006
Nortel Networks Limited
IP Law Group

Date: December 20, 2002